Learning, Hedging, and the Natural Rate Hypothesis[*]

Thomas E. Cone[†] Paul Shea[‡]
SUNY Brockport University of Kentucky

August 6, 2009

Abstract

We assume that firms actively manage risk by learning and hedging in two macroeconomic models, a simple cobweb model and a model of monopolistic competition that allows for the analysis of monetary policy and welfare. In both models, firms that learn (by paying a cost to observe the model's stochastic shocks) face less uncertainty and produce more output than firms that hedge. Parameter or policy changes that increase the attractiveness of learning therefore induce a higher steady state level of output. When we introduce monetary policy into the latter model, we obtain three results that are profoundly different from all or most of the related literature. First, by affecting the volatility of the system, monetary policy affects the supply side decisions of producers and therefore affects the steady state level of output and not just its second moment. Second, if the fraction of firms that learn is exogenous, then the optimal policy maximizes steady state output and welfare by minimizing the volatility of the aggregate price level. Third, if the fraction of firms that learn is endogenous, then the optimal policy maximizes steady state output by ensuring an intermediate level of price stability, close to the minimum stability needed to induce the maximum amount of learning.

JEL Classification:

Keywords:

---

[*] We thank seminar participants at the University of Kentucky for helpful comments. All errors are the authors' responsibility.

[†] E-mail: tcone@brockport.edu

[‡] Email: paul.shea@uky.edu.

1 Introduction

This paper develops a framework where firms use learning and hedging as alternative ways of managing risk. Modern models of the business cycle typically assume that households are risk averse, but that firms are risk neutral. When these models include a role for monetary policy, most notably in the New Keynesian framework, a set of common results emerges. First, although there exists a significant short-term tradeoff between inflation and output, rules based monetary policy has only a very small effect on the natural rate (steady state level) of output. Second, optimal policy should therefore attempt to limit the volatility of output but not attempt to influence its first moment. Third, optimal monetary policy often stabilizes output by stabilizing the aggregate price level.[1]

Although macroeconomic models usually assume risk neutral firms, evidence from the finance literature overwhelmingly suggests that suppliers actively avoid risk. Conventional microeconomic theory suggests that while a "mom and pop" firm should be as risk averse as a typical household, a large and widely held corporation should be close to risk neutral.[2] The Wharton School and Chase Manhattan Bank (1995), however, find that over 70% of corporations hedge against risk through the use of derivatives. Geczy, Minton, and Schrand (1997) find that large corporations also hedge against exchange rate risk through the use of currency derivatives. Theoretical explanations for risk aversion among large corporations include a principal-agent problem involving a risk averse firm manager, the effects of risk on external finance, and heightened costs of financial distress.[3] The principal-agent explanation enjoys significant empirical support. May (1995) examines the risk reduction behavior of American corporations and finds that those with CEOs who have significant wealth vested in the firm engage in more risk reduction. Tufano (1996) studies the gold-mining industry and finds that the degree of firm risk aversion increases with the amount of stock options held by the firm's manager. There is also ample anecdotal evidence of firm risk aversion. In a well documented case, Southwest airlines reaped tremendous financial benefits by having hedged their fuel costs before the market price rose in 2007.[4]

When we introduce this framework into a micro-founded model with rules based monetary policy, the results are surprisingly and profoundly different from all of the related literature. First because aggregate volatility affects the production decisions of firms, and monetary policy affects the stability of the system, monetary policy has supply side effects by inducing firms that face uncertainty to produce more output. These effects influence the steady state level of output and not just its second

---

[1] For details on monetary policy in the New Keynesian literature, see Woodford (2003).

[2] Here "mom and pop" refers to a firm owned by a single household that provides that household's entire income.

[3] See, for example, Amihud and Lev (1981).

[4] See "Southwest Airlines Reaps Benefits of Fuel Hedging Strategy." *Los Angeles Times*, 5/30/08.

moment. Furthermore, the expected utility of the representative household closely tracks the steady state level of output.[5] Second, if the fraction of firms that pay to learn about the model's stochastic shocks is exogenous, then the optimal policy is the rule that stabilizes the aggregate price level. Third, if the choice of whether to learn is endogenous, then excessive price stability leads to less learning and is not optimal. In this case, the optimal policy is to ensure an intermediate level of price stability, equal to the minimum amount needed to cause all firms in the economy to learn. Thus, risk aversion with a choice to learn or hedge can lead to the opposite of the effects that result from risk aversion alone: The pure risk aversion effect is that greater price stability increases mean output. The effect with risk aversion and endogenous learning choice is that greater price stability decreases mean output.

Macroeconomists as diverse as Edward Prescott and Alan Greenspan have suggested in non-academic settings that monetary policy may have real supply side effects by stabilizing the business climate.[6] We are unaware, however, of any previous academic work that formalizes this idea. Other papers have demonstrated that monetary policy may have supply side effects through very different mechanisms. The most common approach is to assume that firms must carry sufficient working capital to pay their wage bill.[7] In these papers, however, the supply side effects of monetary policy on output are limited to its second moment.

The paper is organized as follows. Section 2 introduces our learning versus hedging framework into a simple cobweb model to examine how the two types of risk management interact. As more firms learn and prices become less volatile, the incentives for both additional learning and hedging decline. As a result, a significant fraction of firms will choose not to learn even if hedging is prohibitively expensive. Because hedgers face more uncertainty than learners, the former produce more output than the latter. The steady state level of output is primarily determined by the fraction of firms that choose to learn. Even though, all else equal, larger stochastic shocks or higher hedging costs increase the risk that hedgers face and reduce their output, these changes also induce additional learning and therefore increase steady state output. Section 3 extends the basic model to the case of linear hedging costs. Section 4 then extends our framework into a model which allows for the analysis of monetary policy and welfare. Like the New Keynesian framework, we assume that firms maximize profits under monopolistic competition. Unlike much of that literature, however, we assume that prices are flexible and that non-learning firms must choose their level of output prior to observing the model's stochastic shocks.[8] The major results from the cobweb model carry through. Importantly, we also generate some

---

[5] This result follows Lucas (1987) who demonstrates that for reasonable household utility functions, the welfare effects of fluctuations in output are small when compared to effects on its first moment.

[6] See Edward Prescott, "Five Macroeconomic Myths." *Wall Street Journal*, 12/11/06 and Alan Greenspan's interview on *The Daily Show with John Stewart*, 9/18/07.

[7] See, for example, Christiano, Eichenbaum, and Evans (1997), and Cooley and Nam (1998).

[8] The pricing problem is thus closer to a simple version of the sticky information approach of Mankiw and Reis (2004)

results related to the conduct of monetary policy that stand in stark contrast to the related literature. Because monetary policy affects the stability of the model, and the stability of the model affects firms' production choices, policy has supply side effects on the first moment of output. For a fixed fraction of learning firms, monetary policy maximizes the expected utility of the representative household by stabilizing prices which induces hedging firms to increase their production. If the fraction of learning firms is endogenous, however, then excessive price stability induces too many firms to hedge which results in lower output. Optimal policy is thus to ensure the minimum amount of price volatility needed to induce all firms to learn instead of hedge. Section 5 provides microeconomic foundations for hedging and learning by assuming that firms must employ labor to perform each activity. If the cost of learning is sufficiently low, then optimal policy continues to ensure enough price volatility to maximize learning. If the costs of learning are sufficiently high, however, then optimal policy stabilizes prices and discourages learning. Finally, Section 6 concludes.

2 The Basic Model

2.1 The Environment

Firms in a market are exposed to risk because their output price is uncertain at the time they choose their output levels. Firms are risk-averse, and their risk aversion is embodied in a mean-variance utility function over profit. They have two alternatives for managing risk; they can either hedge or learn. They can also do neither, but not both.

In this basic model firms are price takers; a learning firm's revenue is $P_tY_{l,t}$ where P_t is the date t output price and $Y_{l,t}$ is the learning firm's date t output. The default information set I_{t-1} includes knowledge of the model, the current fraction of learning firms $q_{l,t}$, and all variables through the end of period $t-1$. To learn, a firm pays cost $\psi \geq 0$ and augments its information set to I_t by observing a demand shock U_t (see the next section).

P^s_{t-1} is the strike price in a hedging position entered into at the end of $t-1$, essentially a bet that P_t will be below P^s_{t-1}. The gain (ignoring transaction costs) from one hedging contract is $P^s_{t-1} - P_t$. A hedging firm produces $Y_{h,t}$ units of output and enters into Z_t units of the hedging contract. Hedging entails transaction cost hZ^a_t and we refer to $h \geq 0$ as the hedging cost. We study quadratic hedging costs ($a = 2$) and linear hedging costs ($a = 1$).[9]

The output cost function is $Y^2_{i,t}$, $i = l, h$. The general optimization problem, involving the choice

---

than the staggered pricing assumption used in much of the New Keynesian literature.

[9] A firm could choose to neither learn nor hedge by "hedging" with $Z_t = 0$, although this never happens in equilibrium with $a = 2$. It does happen in some equilibria with $a = 1$.

variables $Y_{l,t}$ for a learning firm and $Y_{h,t}$ and Z_t for a hedging firm, is

$$\max \Pi_{i,t} = E(F_{i,t}|I_t^*) - \phi var(F_{i,t}|I_t^*) - f_{i,t} \tag{1}$$

where $F_{i,t}$ is profit from output sales, $f_{i,t}$ is the cost of learning or hedging, ϕ is a risk aversion parameter, and I_t^* is the relevant information set, $I_t^* = I_t$ for a learning firm and $I_t^* = I_{t-1}$ for a hedging firm. To make the hedging-versus-learning choice, a firm solves (1) under each alternative and then compares their expected utilities conditional on I_{t-1} . Nash equilibrium among firms determines $q_{l,t}$ and equilibrium is non-trivial because $q_{l,t}$ affects the expectation and variance of P_t conditional on I_{t-1} and these moments in turn affect the set of equilibrium $q_{l,t}$'s.

We first specialize $a = 2$ (the linear hedging cost case $a = 1$ is treated below). A hedger's revenue is $(Y_t - Z_t)P_t + Z_t P_{t-1}^s$ with mean and variance $(Y_t - Z_t)E(P_t|I_{t-1}) + Z_t P_{t-1}^s$ and $(Y_t - Z_t)^2 var(P_t|I_{t-1})$ respectively, so a hedger's optimization problem is

$$\max_{Y_{h,t}, Z_t} \Pi_{h,t} = (Y_{h,t} - Z_t)E(P_t|I_{t-1}) + Z_t P_{t-1}^s - Y_{h,t}^2 - \phi(Y_{h,t} - Z_t)^2 var(P_t|I_{t-1}) - hZ_t^2. \tag{2}$$

We immediately impose $P_{t-1}^s = E(P_t|I_{t-1})$ and obtain the first-order conditions[10]

$$Y_{h,t} = \frac{E(P_t|I_{t-1})}{2} \frac{h + \phi var(P_t|I_{t-1})}{h + (1 + h)\phi var(P_t|I_{t-1})} \tag{3}$$

$$Z_t = \frac{E(P_t|I_{t-1})}{2} \frac{\phi var(P_t|I_{t-1})}{h + (1 + h)\phi var(P_t|I_{t-1})}. \tag{4}$$

Note the fraction of output hedged

$$\frac{Z_t}{Y_{h,t}} = \frac{\phi var(P_t|I_{t-1})}{h + \phi \; var(P_t|I_{t-1})}$$

is less than one for positive h, and $\frac{\partial Z_t / Y_{h,t}}{\partial h} < 0$. Also, $Y_{h,t}$ is strictly decreasing as a function of $var(P_t|I_{t-1})$. Substitute (3) and (4) into (2) to find the utility conditional on I_{t-1}:

$$\Pi_{h,t}^* = \frac{[E(P_t|I_{t-1})]^2}{4} \frac{h + \phi var(P_t|I_{t-1})}{h + (1 + h)\phi var(P_t|I_{t-1})}. \tag{5}$$

To assess the desirability of learning a firm solves the learner's problem conditional on I_t, calculates the resultant utility, then takes the expectation of that utility conditional on I_{t-1}. A learner's optimization problem is

$$\begin{aligned} \max_{Y_t} \Pi_{l,t} &= Y_{l,t}E(P_t|I_t) - Y_{l,t}^2 - \phi Y_{l,t}^2 var(P_t|I_t) - \psi \qquad (6) \\ &= Y_{l,t}P_t - Y_{l,t}^2 - \psi \end{aligned}$$

---

[10] This avoids having the price in the hedging market reveal some information about P_t: If $P_{t-1}^s = E(P_t|I_t)$, observing P_{t-1}^s does not affect firms' expectations of P_t.

where $E(P_t|I_t) = P_t$ and $var(P_t|I_t) = 0$ because a learner observes the demand shock. The first-order condition is

$$Y_{l,t} = \frac{P_t}{2}. \tag{7}$$

Note learners' expected output is always higher than hedgers' expected output for $h > 0$, because

$$E(Y_{h,t}|I_{t-1}) = Y_{h,t} = \frac{E(P_t|I_{t-1})}{2} \frac{h + \phi var(P_t|I_{t-1})}{h + (1+h)\phi var(P_t|I_{t-1})}$$

and

$$E(Y_{l,t}|I_{t-1}) = \frac{E(P_t|I_{t-1})}{2}.$$

This fact, which is also true in the more sophisticated model of Section 4, has important implications for monetary policy.

Equation (7) implies utility

$$\frac{P_t^2}{4} - \psi, \tag{8}$$

so a firm evaluates the attractiveness of learning *ex ante* as

$$\Pi^*_{l,t} = \frac{E\left[P_t^2|I_{t-1}\right]}{4} - \frac{\phi var(P_t^2|I_{t-1})}{16} - \psi. \tag{9}$$

To choose between learning and hedging firms compare (5) with (9). Because the expectations and variances in those expressions are functions of $q_{l,t}$, some regions of the parameter space have multiple equilibrium values of $q_{l,t}$, as we document below. An equilibrium with all firms learning exists if

$$\Pi^*_{h,t} \leq \Pi^*_{l,t} \tag{10}$$

holds for $q_{l,t} = 1$. An equilibrium with no firms learning exists if

$$\Pi^*_{h,t} \geq \Pi^*_{l,t} \tag{11}$$

holds for $q_{l,t} = 0$. An equilibrium with learning and hedging exists if

$$\Pi^*_{h,t} = \Pi^*_{l,t} \tag{12}$$

holds for one or more values of $q_{l,t} \in (0, 1)$.

Two of the factors that influence firms' learn-or-hedge choice are the risk aversion effect and the Jensen's Inequality effect. The risk aversion effect is embodied in the variance part of the mean-variance objective function and—roughly speaking—makes firms want to hedge. The "Jensen's Inequality effect" refers to the convexity of the profit function in the price, which is embodied in the mean part of the objective function. Firms can only take advantage of this convexity if they know the price, so this effect provides incentive for firms to learn. This is brought out most clearly if hedging and learning are costless. In that case the comparison is between

$$\Pi^*_{h,t} = \frac{[E(P_t|I_{t-1})]^2}{4}$$

and

$$\Pi^*_{l,t} = \frac{E\left[P_t^2|I_{t-1}\right]}{4} - \frac{\phi var(P_t^2|I_{t-1})}{16}.$$

Jensen's Inequality implies $E\left[P_t^2|I_{t-1}\right] > [E(P_t|I_{t-1})]^2$ so that a firm would always learn if the term $-\frac{\phi var(P_t^2|I_{t-1})}{16}$ were not present in $\Pi^*_{l,t}$. That term captures the risk of learning.

2.2 Equilibrium

We have six endogenous variables for which to solve: $Y_{l,t}$, $Y_{h,t}$, Z_t, P_t, $var(P_t|I_{t-1})$, and $q_{l,t}$ (hereafter we omit the time subscript on $q_{l,t}$). We must solve for $var(P_t|I_{t-1})$ because it affects firms' learning versus hedging choice and the amount of output that hedgers produce. To solve for these variables we have six relevant equations or inequalities; hedgers' two first-order conditions (3) and (4), learners' first-order condition (7), the demand equation (14) (see below), the definition of the variance, and the equation or inequality that compares hedgers' expected utility to learners', (10) - (12).

Using (3) and (7), total supply is

$$Y_t = \left(\frac{P_t}{2}\right)^{q_l}\left(\frac{E(P_t|I_{t-1})}{2}\frac{h+\phi var(P_t|I_{t-1})}{h+(1+h)\phi var(P_t|I_{t-1})}\right)^{1-q_l}. \tag{13}$$

Demand is

$$D_t = \frac{M_I U_t}{P_t^{m_p}} \tag{14}$$

where M_I and m_p are positive parameters and U_t is a positive shock uniformly distributed on $[\rho, \omega]$. Market clearing gives

$$P_t = (2M_I)^{\frac{1}{m_p+1}}\left[E(U_t^{\frac{1}{m_p+q_l}}|I_{t-1})\right]^{\frac{q_l-1}{m_p+1}}\left(\frac{h+(1+h)\phi var(P_t|I_{t-1})}{h+\phi var(P_t|I_{t-1})}\right)^{\frac{1-q_l}{m_p+1}} U_t^{\frac{1}{m_p+q_l}} \tag{15}$$

which implies

$$var(P_t|I_{t-1}) = (2M_I)^{\frac{2}{m_p+1}}\left(\frac{h+(1+h)\phi var(P_t|I_{t-1})}{h+\phi var(P_t|I_{t-1})}\right)^{\frac{2(1-q_l)}{m_p+1}}\frac{var(U_t^{\frac{1}{m_p+q_l}}|I_{t-1})}{\left[E(U_t^{\frac{1}{m_p+q_l}}|I_{t-1})\right]^{\frac{2(1-q_l)}{m_p+1}}}. \tag{16}$$

Equations (15) and (16) are not solutions for P_t and $var(P_t|I_{t-1})$ because they are expressed as functions of q_l. Furthermore, the equation for the variance cannot be solved analytically even for known q_l, so the system cannot be solved analytically, in general. The system can be solved computationally for specific parameters. Before proceeding to that we present some analytical results.

2.3 Analytical Results

Arbitrarily high hedging costs do not necessarily induce all firms to learn.

Proposition 1. *In a (nonempty) region of the parameter space, equilibria with* $q_l = 1$ *do not survive as hedging cost* h *goes to infinity.*

Proof. The inequality that determines the learning-hedging choice is

$$\frac{[E(P_t|I_{t-1})]^2}{4}\frac{h+\phi var(P_t|I_{t-1})}{h+(1+h)\phi var(P_t|I_{t-1})} \lesseqgtr \frac{E\left[P_t^2|I_{t-1}\right]}{4} - \frac{\phi var(P_t^2|I_{t-1})}{16} - \psi. \tag{17}$$

If the left side of (17) is greater than the right side, hedging is preferred to learning. Substituting in the moments $E(P_t|I_{t-1})$, $var(P_t|I_{t-1})$, $[P_t^2|I_{t-1}]$, $var(P_t^2|I_{t-1})$ and rearranging gives

$$X_t^2 \phi var(U_t^{\frac{2}{m_p+q_l}}|I_{t-1}) - X_t \left\{ E\left[U_t^{\frac{2}{m_p+q_l}}|I_{t-1}\right] - \frac{h+\phi var(P_t|I_{t-1})}{h+(1+h)\phi var(P_t|I_{t-1})}[E(U_t^{\frac{1}{m_p+q_l}}|I_{t-1})]^2 \right\} \lesseqgtr -\psi$$

where

$$X_t \equiv \frac{(2M_I)^{\frac{2}{m_p+1}}}{4}\left(\frac{\frac{h+(1+h)\phi var(P_t|I_{t-1})}{h+\phi var(P_t|I_{t-1})}}{E(U_t^{\frac{1}{m_p+q_l}}|I_{t-1})}\right)^{\frac{2(1-q_l)}{m_p+1}}.$$

L'Hopital's Rule implies $\lim_{h\to\infty}\frac{h+(1+h)\phi var(P_t|I_{t-1})}{h+\phi var(P_t|I_{t-1})} = 1+\phi var(P_t|I_{t-1})$, so as $h \to \infty$ the inequality becomes

$$(\lim_{h\to\infty} X_t)^2 \phi var(U_t^{\frac{2}{m_p+q_l}}|I_{t-1}) - (\lim_{h\to\infty} X_t)\left[E\left(U_t^{\frac{2}{m_p+q_l}}|I_{t-1}\right) - \frac{[E(U_t^{\frac{1}{m_p+q_l}}|I_{t-1})]^2}{1+\phi var(P_t|I_{t-1})}\right] \lesseqgtr -\psi. \tag{18}$$

where

$$\lim_{h\to\infty} X_t = \frac{(2M_I)^{\frac{2}{m_p+1}}}{4}\left(\frac{1+\phi var(P_t|I_{t-1})}{E(U_t^{\frac{1}{m_p+q_l}}|I_{t-1})}\right)^{\frac{2(1-q_l)}{m_p+1}}.$$

Next we show that $q_l = 1$ implies a contradiction. Imposing $q_l = 1$ gives

$$\frac{(2M_I)^{\frac{4}{m_p+1}}}{16}\phi var(U_t^{\frac{2}{m_p+1}}|I_{t-1}) - \frac{(2M_I)^{\frac{2}{m_p+1}}}{4}\left[E(U_t^{\frac{2}{m_p+1}}|I_{t-1}) - \frac{[E(U_t^{\frac{1}{m_p+1}}|I_{t-1})]^2}{1+\phi var(P_t|I_{t-1})}\right] \lesseqgtr -\psi. \tag{19}$$

The only endogenous variable in (19) is $var(P_t|I_{t-1})$. Also, the moments of U_t are finite and the term in brackets is bounded in $[var(U_t^{\frac{1}{m_p+1}}|I_{t-1}), E(U_t^{\frac{2}{m_p+1}}|I_{t-1})]$ as $var(P_t|I_{t-1})$ varies in $\Re_+$. It follows that for sufficiently large ψ the left side is larger than the right side. Therefore, by contradiction, for large enough ψ equilibria with $q_l = 1$ do not survive as $h \to \infty$. ■

Not only is an arbitrarily high hedging cost consistent with some firms not learning, but it is consistent with no firms learning, as the next proposition attests:

Proposition 2. *For a (nonempty) region of the parameter space, as $h \to \infty$ the only equilibrium value of q_l is zero.*

Proof. The previous proposition implies that we only need consider $q_l < 1$. Send $\psi \to \infty$. If $var(P_t|I_{t-1})$ is bounded as $\psi \to \infty$ then the left side of (18) goes to a finite (possibly negative) limit. Then the left side is larger than the right side, so hedging is preferred to learning. If $var(P_t|I_{t-1}) \to \infty$ as $\psi \to \infty$ then the left side of (18) goes to (positive) infinity, so the left side is larger than the right side, so hedging is preferred to learning. ■

While arbitrarily high hedging costs do not necessarily eliminate "hedging," arbitrarily high learning costs do eliminate learning:

Proposition 3. *For sufficiently high learning cost ψ, the only equilibrium value of q_l is zero.*

Proof. Rearranging the inequality that determines the learning-hedging choice gives

$$\frac{[E(P_t|I_{t-1})]^2}{4}\frac{h+\phi var(P_t|I_{t-1})}{h+(1+h)\phi var(P_t|I_{t-1})}+\frac{\phi var(P_t^2|I_{t-1})}{16} \lessgtr \frac{E\left[P_t^2|I_{t-1}\right]}{4}-\psi. \tag{20}$$

If the left side of (20) is greater than the right side, hedging is preferred to learning. The left side is strictly positive, so showing that the right side is negative for sufficiently large ψ suffices to prove the proposition. To show that, in turn, it suffices to show that $E\left[P_t^2|I_{t-1}\right]$ is bounded. From (15) we have

$$E\left[P_t^2|I_{t-1}\right]=(2M_I)^{\frac{2}{m_p+1}}\left[E(U_t^{\frac{1}{m_p+q_l}}|I_{t-1})\right]^{\frac{2[q_l-1]}{m_p+1}}\left(\frac{h+(1+h)\phi var(P_t|I_{t-1})}{h+\phi var(P_t|I_{t-1})}\right)^{\frac{2[1-q_l]}{m_p+1}}E(U_t^{\frac{2}{m_p+q_l}}|I_{t-1}) \tag{21}$$

and the right side of (21) is bounded as ψ varies: The only variables on the right side that are potentially functions of ψ are $var(P_t|I_{t-1})$ and q_l. As they vary within their admissible ranges ($\Re_+$ and $[0,1]$ respectively) the RHS remains within a bounded interval. It follows that the right side of (20) goes to negative infinity as $\psi \to \infty$. ■

It should be noted that the functional forms for the learning and hedging costs drive the asymmetry of these results. The hedging cost is continuous, so firms can hedge and manage their total hedging costs, as h becomes large, by hedging a small amount. This is important for Propositions 1 and 2. In contrast, the learning cost is binary (an assumption made for tractability), so learning firms cannot manage their total learning costs as ψ becomes large. This is important for Proposition 3.

2.4 Simulation Results

Next we parameterize and simulate the model. Throughout we maintain $a = 2$, $M_I = 1$, $m_p = 1.1$. We hold $\rho = 0.01$ and $\omega = 1.99$ unless noted otherwise. We study different values of the risk aversion ϕ, the learning cost ψ, and the hedging cost h. In searching for equilibria for a given parameterization, our simulation searches over $q_l \in \{0, 0.01, 0.02, ..., 1\}$.[11]

2.4.1 No learning or hedging costs, $h = \psi = 0$

In simulations with $h = \psi = 0$ sufficiently low risk aversion ϕ (less than 0.52) entails universal learning as the only equilibrium—this is an example of the Jensen's Inequality effect as explained in section 2.1—while sufficiently high ϕ (greater than 0.68) entails universal hedging as the only equilibrium. The intermediate ranges of ϕ have multiple equilibrium values of q_l.

For example, with $\phi = 0.65$ there is an unstable equilibrium at $q_l = 0.67$ and stable equilibria at $q_l = 0$ and $q_l = 1$ (see Figure 4). The dual stable corner equilibria exist due to the way that q_l affects $\text{var}(P_t|I_{t-1})$: A high q_l means most firms know the demand shock. That makes supply quite responsive to the demand shock, which dampens price fluctuations, i.e., lowers $\text{var}(P_t|I_{t-1})$. Because the price variance is low, firms face little risk, and so are willing to remain unhedged. Thus an individual firm's best response is to learn, and the (notional time) best response dynamics push the economy to $q_l = 1$. Symmetrically, for low values of q_l, the price variance is high and the high risk makes firms prefer to hedge, driving the economy to $q_l = 0$. The unstable interior equilibrium is of course at the point where the direction of these effects changes.[12]

2.4.2 No hedging costs, $h = 0$, $\psi > 0$

The set of equilibria is quite sensitive to the learning cost. With the above parameterization, we choose $\phi = 0.5$—a level of risk aversion for which universal learning is the sole equilibrium if $h = \psi = 0$ (see section 2.4.1)—and $\psi = 0.01$. With this small learning cost, stable equilibria exist at each corner and an unstable equilibrium at $q_l = 0.11$ (see Figure 4). Raising ψ to 0.015 does two things: First, it raises the unstable equilibrium to $q_l = 0.39$, that is, it enlarges the interval in which the best response dynamics push the economy to zero learning. Second, it lowers the upper stable equilibrium from $q_l = 1$ to $q_l = 0.61$, and in this sense also discourages learning. Raising ψ further to 0.02 leaves only the $q_l = 0$ equilibrium.

Maintaining ψ at 0.02, lowering ϕ restores the $q_l = 0$ and the interior equilibria.

Even with zero hedging cost and positive learning cost, universal learning can be the only equilibrium due to the Jensen's Inequality effect. An example parameterization is $\phi = 0.05$, $h = 0$, $\psi = 0.01$.

---

[11] When an equilibrium is located between two values for q_l we report the higher value.

[12] In our simulations, unsurprisingly for an unstable equilibrium, this equilibrium has counterintuitive comparative statics: Lower risk aversion leads to less learning and more hedging.

Figure 1: Firms' utility for $h = \psi = 0$, $\varphi = 0.65$.

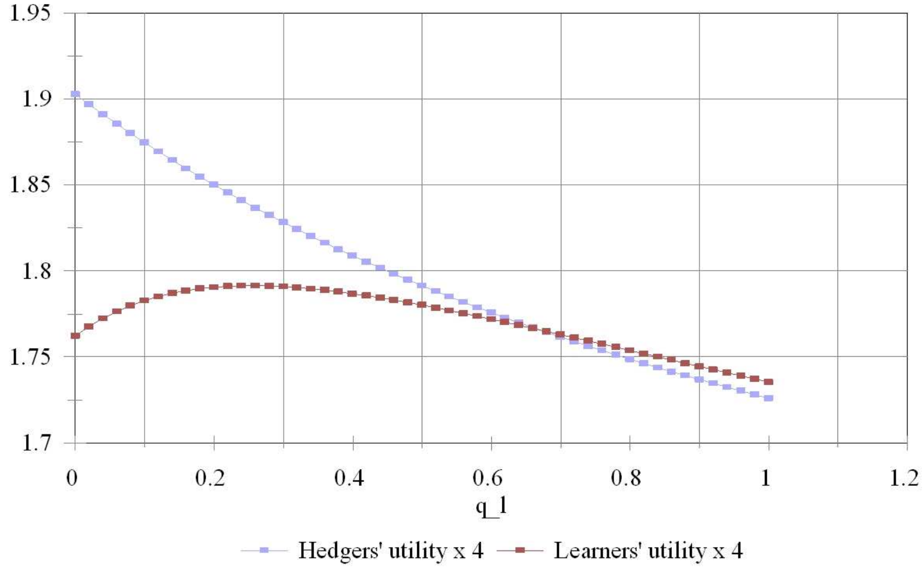

2.4.3 No learning costs, $\psi = 0$, $h > 0$

We next raise h. We also raise ϕ to 0.7 so that we start, with $h = 0$, from a parameterization for which the only equilibrium is universal hedging (see section 2.4.1). Raising the hedging cost to $h = 0.05$ creates dual stable corner equilibria and an unstable equilibrium at $q_l = 0.34$. Further increases in h push the interior equilibrium toward $q_l = 0$, until it disappears and the only equilibrium has $q_l = 1$.

2.4.4 Positive hedging and learning costs

Positive h and ψ are consistent with an equilibrium value of q_l at either endpoint or in the interior of $[0, 1]$, as well as multiple equilibria.

A typical parameterization of the ones we study has $\phi = 0.5$, $h = 1$, and $\psi = 0.1$. This economy has one equilibrium, $q_l = 0$. Studying other values of q_l (i.e., as if q_l were exogenous), we see that as q_l rises, the expectation and variance of the price fall. The expected price falls because learners produce more in expectation than hedgers, and the price variance falls because learners' output responds to demand shocks, thus dampening price fluctuations. For these same reasons, the variance of output

Figure 2: Firms' utility for $h = 0$, $\psi = 0.01$ and $\psi = 0.015$, $\varphi = 0.5$.

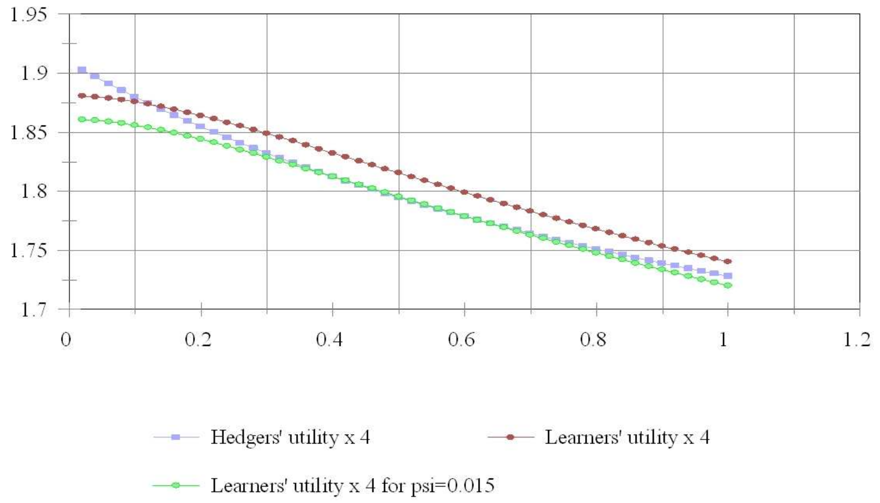

rise as q_l rises, and expected output typically rises.[13] This behavior of the endogenous variables is typical in our simulations.

We can produce some learning by reducing firms' risk aversion or the learning cost. Reducing firms' risk aversion to $\phi = 0.25$ implies a unique equilibrium, which is stable, at $q_l = 0.25$. Reducing the learning cost to $\psi = 0.05$ implies a unique equilibrium, which is stable, at $q_l = 0.94$.

2.4.5 Neither Hedging Nor Learning

Firms can choose to neither hedge nor learn, a choice encompassed by choosing "hedging" with $Z_t = 0$. The possibility of doing neither leads to some system behavior that seems strange; in particular driving the hedging cost to infinity does not necessarily drive q_l to unity, as attested to in Propositions 1 and 2. This occurs because firms that "hedge" are effectively hedging zero units, so increases in the hedging cost do not affect their behavior. Note, though, that because the hedging cost is Inada-convex, hedging firms will never hedge *exactly* zero units.

---

[13] For q_l in $[0, 0.05]$, higher q_l is associated with lower expected output. In that range of q_l, the fact that hedgers and learners produce less output as q_l rises, counteracts the fact that more firms are becoming learners, who produce more than hedgers.

The following parameterizations illustrate this phenomenon. With $\phi = 1$, $\psi = 0.05$, and $h = 2.5$, the only equilibrium exists at $q_l = 0$ and it is stable. For higher hedging cost $h = 5$ the $q_l = 0$ equilibrium persists, and there is also an unstable equilibrium at $q_l = 0.4$ and a stable equilibrium at $q_l = 0.82$. Continuing to raise h moves the unstable equilibrium downward and the stable equilibrium upward, until they reach 0.32 and 0.92 respectively for $h = 40$. Beyond that point, increases in h have no discernible effect on the set of equilibria. Studying the amount of output hedged reveals why: "Hedging" firms essentially hedge zero units. Since "hedging" firms are hedging essentially zero units, further increases in the hedging cost do not affect their behavior.[14]

Why don't all firms learn as $h \to \infty$? The learning cost is too high to be worth it. This is due to the fact that the learning cost is binary (a firm learns or it doesn't learn, an assumption made for tractability) while the hedging cost is continuous. No matter how high h is, firms can hedge an arbitrarily small amount and thus pay an arbitrarily small amount. On the other hand, the fixed learning cost ψ implies that the benefit of learning is not worth the cost.

2.4.6 Exogenous Volatility, Endogenous Volatility, and Output

To study how the learning-hedging choice affects the economy's volatility, we examine different levels of the exogenous shock's variance. (We also raise the shock's mean so we have more scope to change its variance.) There are two main results of interest: One is that a higher variance of the shock can cause a reduction in the variance of the price. The other is that higher exogenous variance can increase mean output, which is the opposite of the effect of risk aversion without the learn-hedge choice.[15] Table 1 presents results for the following parameters, which we selected because they entail stable interior equilibria in q_l: For $M_I = 1$, $m_p = 1.1$, $\phi = 0.1$, $\psi = 0.1$, $h = 1$, $\rho = 6$, and $\omega = 14$, the moments of U are $E(U|I_{t-1}) = 10$ and $\text{var}(U|I_{t-1}) = 5.333$. We also apply a mean-preserving spread to U by setting $\rho = 5$, $\omega = 15$, which implies $\text{var}(U|I_{t-1}) = 8.333$. This induces more firms to learn, and because learning firms respond to the demand shock, amplifies output volatility and dampens price volatility. It also raises mean output and lowers the mean price because learners produce more (in expectation) than hedgers. (The higher mean output is not due to firms within each group producing more; each group produces less with the high-variance shock than it does with the low-variance shock.)

---

[14] Even for $h = 1,000,000$ the set of equilibria remains the same as for $h = 40$.

[15] If firms were risk averse and could not learn or hedge, then their output would be

$$Y_t = \frac{E(P_t|I_t)}{2\left[1 + \phi \text{var}(P_t|I_t)\right]},$$

which is decreasing in $\text{var}(P_t|I_t)$. Intuitively, if the market is riskier, risk-averse firms want a smaller stake in the market.

Table 1

var$(U|I_{t-1})$ **and Endogenous Variables**

($M_I = 1$, $m_p = 1.1$, $\phi = 0.1$, $\psi = 0.1$, $h = 1$)

var$(U\|I_{t-1})$	q_l	$E(Y\|I_{t-1})$	var$(Y\|I_{t-1})$	$E(P\|I_{t-1})$	var$(P\|I_{t-1})$
5.3333	0.35	2.0373	0.0137	4.2030	0.4601
8.3333	0.75	2.0459	0.0621	4.1449	0.4450

When we decrease the variance still further all firms hedge, and when we increase the variance still further all firms learn. This illustrates the Jensen's Inequality effect, which encourages learning. That is, the larger is the demand shock's variance, the more the convexity of profit in the price matters, so the stronger is the incentive to learn. However, in simulations with higher risk aversion ϕ, the risk aversion effect dominates, and a higher variance discourages learning instead of encouraging it.

Note that in the foregoing simulations, the fundamental risk as measured by var$(U|I_{t-1})$ and the risk aversion do not necessarily have similar effects: Raising the risk aversion parameter always makes learning less attractive compared to hedging, while raising var(U) may make learning more or less attractive, depending on which effect dominates: The Jensen's Inequality effect in the first part of the objective function, or the risk aversion effect in the second part of the objective function. And of course, firms' equilibrium adjustments in q_l complicate things more, e.g., as Table 1 attests, they can make var$(U|I_{t-1})$ and var$(P|I_{t-1})$ move in opposite directions.

3 Linear Hedging Cost

With $a = 1$ the hedging cost is the linear function hZ_t. Because non-learners hedge exactly zero units in some equilibria (in contrast to the quadratic cost case), we distinguish among the three possibilities—learning, hedging, and doing neither—in this section. Conveniently, it can be shown analytically that there will never be positive measure of firms choosing both learning and hedging in the same equilibrium.

We impose $Y_{h,t}, Z_t \geq 0$ and it can be shown that possible solutions to a non-learner's problem have $0 \leq Z_t \leq Y_{h,t}$ with at least one inequality strict. The following work pertains to the possible solutions; we omit the demonstration that these possibilities are exhaustive. The problem in Lagrangian form is

$$\max_{\lambda_Z,\lambda_Y,Y_{h,t},Z_t} L_{h,t} = (Y_{h,t}-Z_t)E(P_t|I_{t-1})+Z_tP^s_{t-1}-Y^2_{h,t}-\phi(Y_{h,t}-Z_t)^2\text{var}(P_t|I_{t-1})-hZ_t+\lambda_Z Z_t+\lambda_Y Y_{h,t}.$$

FOCs:

$$\lambda_Y, \lambda_Z \geq 0$$

$$Y_{h,t}, Z_t \geq 0$$

$$\lambda_Y Y_{h,t} = \lambda_Z Z_t = 0$$

$$\frac{\partial L_{h,t}}{\partial Y_{h,t}} = E(P_t|I_{t-1}) - 2Y_{h,t} - 2\phi \text{var}(P_t|I_{t-1})(Y_{h,t} - Z_t) + \lambda_Y = 0$$

$$\frac{\partial L_{h,t}}{\partial Z_t} = 2\phi \text{var}(P_t|I_{t-1})(Y_{h,t} - Z_t) - h + \lambda_Z = 0.$$

(1) Suppose the solution has $Y_{h,t}$ and Z_t positive. Then the multipliers must be zero, so we have

$$Y_{h,t} = \frac{E(P_t|I_{t-1}) - h}{2}$$

$$Z_t = \frac{E(P_t|I_{t-1}) - h}{2} - \frac{h}{2\phi \text{var}(P_t|I_{t-1})}. \tag{22}$$

Optimized *ex ante* utility for a hedger is

$$\Pi^*_{h,t} = \frac{[E(P_t|I_{t-1}) - h]^2}{4} + \frac{h^2}{4\phi \text{var}(P_t|I_{t-1})}. \tag{23}$$

(2) Suppose $Z_t = 0 < Y_{h,t}$. Then $\lambda_Y = 0$ so we have

$$Y_{n,t} = \frac{E(P_t|I_{t-1})}{2[1 + \phi textnormalvar(P_t|I_{t-1})]}$$

$$\lambda_Z = h - \frac{\phi \text{var}(P_t|I_{t-1}) E(P_t|I_{t-1})}{1 + \phi \text{var}(P_t|I_{t-1})}$$

where the n subscript stands for "nothing." This solution works provided the RHS of the λ_Z equation is non-negative. In this case the optimized utility is

$$\Pi^*_{n,t} = \frac{[E(P_t|I_{t-1})]^2}{4[1 + \phi \text{var}(P_t|I_{t-1})]}. \tag{24}$$

This cannot be a solution when the $Y, Z > 0$ possibility is a solution. That is, it is not possible for positive measure of non-learning firms to be hedging and positive measure to be not hedging in the same equilibrium. The proof comes from rearranging the FOC for Z in that case, (22):

$$Z_t = \frac{E(P_t|I_{t-1}) - h}{2} - \frac{h}{2\phi \text{var}(P_t|I_{t-1})}.$$

Since the $Z > 0$ requires this to be positive, we have

$$\frac{\phi \text{var}(P_t|I_{t-1}) E(P_t|I_{t-1})}{1 + \phi \text{var}(P_t|I_{t-1})} > h$$

which directly contradicts the requirement that $\lambda_Z \geq 0$ in the $Z_t = 0 < Y_{h,t}$ case.

Notice that learners produce more in expectation than hedging firms and firms that do nothing, a result that survives the change in the hedging cost function. One difference is that if a firm hedges, its output is not a function of the price variance.

3.1 Equilibrium

3.1.1 Case 1: All non-learners hedge

If all non-learners hedge, total supply is

$$Y_t = \left(\frac{P_t}{2}\right)^{q_l} \left(\frac{E(P_t|I_{t-1}) - h}{2}\right)^{1-q_l}.$$

Demand:

$$D_t = \frac{M_I U_t}{P_t^{m_p}}$$

Market clearing:

$$P_t = \left(\frac{2M_I U_t}{(E(P_t|I_{t-1}) - h)^{1-q_l}}\right)^{\frac{1}{m_p+q_l}}.$$

Take expectations:

$$E(P_t|I_{t-1}) = \left(\frac{2M_I}{(E(P_t|I_{t-1}) - h)^{1-q_l}}\right)^{\frac{1}{m_p+q_l}} E(U_t^{\frac{1}{m_p+q_l}}|I_{t-1}).$$

The variance:

$$\text{var}(P_t|I_{t-1}) = \left(\frac{2M_I}{(E(P_t|I_{t-1}) - h)^{1-q_l}}\right)^{\frac{2}{m_p+q_l}} \text{var}(U_t^{\frac{1}{m_p+q_l}}|I_{t-1}).$$

The expectation and variance of output:

$$E(Y_t|I_{t-1}) = \frac{1}{2}E(P_t^{q_l}|I_{t-1}) \left[E(P_t|I_{t-1}) - h\right]^{1-q_l}$$

$$\text{var}(Y_t|I_{t-1}) = \frac{\left[E(P_t|I_{t-1}) - h\right]^{2(1-q_l)}}{4}\text{var}(P_t^{q_l}|I_{t-1}).$$

3.1.2 Case 2: If no firm hedges

If no firm hedges, all firms either learning or doing nothing, total supply is

$$Y_t = \left(\frac{P_t}{2}\right)^{q_l} \left(\frac{E(P_t|I_{t-1})}{2[1 + \phi\text{var}(P_t|I_{t-1})]}\right)^{1-q_l}.$$

Market clearing:

$$P_t = (2M_I)^{\frac{1}{m_p+1}} U_t^{\frac{1}{m_p+q_l}} \left(\frac{1 + \phi\text{var}(P_t|I_{t-1})}{E(U_t^{\frac{1}{m_p+q_l}}|I_{t-1})}\right)^{\frac{1-q_l}{m_p+1}}$$

$$E(P_t|I_{t-1}) = (2M_I)^{\frac{1}{m_p+1}} \left[E(U_t^{\frac{1}{m_p+q_l}}|I_{t-1})\right]^{\frac{m_p+q_l}{m_p+1}} (1 + \phi\text{var}(P_t|I_{t-1}))^{\frac{1-q_l}{m_p+1}}$$

The variance:

$$\mathrm{var}(P_t|I_{t-1}) = (2M_I)^{\frac{2}{1+m_p}} \left(\frac{1+\phi \mathrm{var}(P_t|I_{t-1})}{E(U_t^{\frac{1}{m_p+q_l}}|I_{t-1})} \right)^{\frac{2(1-q_l)}{1+m_p}} \mathrm{var}(U_t^{\frac{1}{m_p+q_l}}|I_{t-1}).$$

3.2 Simulation Results

We parameterize as in the quadratic case. Throughout we maintain $M_I = 1$, $m_p = 1.1$, and we use $\rho = 0.01$ and $\omega = 1.99$ unless noted otherwise. We study different values of ϕ, ψ, and h and search for equilibria over q_l.

3.2.1 No hedging costs, $h = 0$, $\psi \geq 0$

With $h = 0$ the value of a is irrelevant, so the results are the same as in sections 2.4.1 and 2.4.2.

3.2.2 No learning costs, $\psi = 0$, $h > 0$

We next raise h. We also raise ϕ to 0.7 so that we start, with $h = 0$, from a parameterization for which the only equilibrium is universal hedging (see section 2.4.1). Raising the hedging cost to $h = 0.05$ creates dual stable equilibria at $q_l = 0$ and $q_l = 1$ and an unstable equilibrium at $q_l = 0.17$, where the alternative to learning is hedging. Futher increases in h push the interior equilibrium toward $q_l = 0$, until for high enough h the only equilibrium has $q_l = 1$.

3.2.3 Positive hedging and learning costs

Positive h and ψ are consistent with an equilibrium value of q_l at either endpoint or in the interior of $[0, 1]$, as well as multiple equilibria. A typical parameterization has $\phi = 0.5$, $h = 1$, and $\psi = 0.1$. This economy has a stable equilibrium at $q_l = 0.25$, with non-learners doing nothing. Not surprisingly, we can induce all firms to learn or hedge by reducing ψ or h, or induce all firms to do neither by setting ψ and h high enough.

Multiple equilibria exist with the following parameters: $\phi = 0.2$, $\psi = 0.05$, $h = 0.2$, $\rho = 1$, $\omega = 11$. Each corner is a stable equilibrium and $q_l = 0.69$ is an unstable equilibrium, with non-learners hedging.

3.2.4 Exogenous Volatility, Endogenous Volatility, and Output

Table 2 presents results for the parameters used in section 2.4.6: $M_I = 1$, $m_p = 1.1$, $\phi = 0.1$, $\psi = 0.1$, $h = 1$, $\rho = 6$, and $\omega = 14$, for which the moments of U are $E(U|I_{t-1}) = 10$ and var$(U|I_{t-1}) = 5.333$. We also apply a mean-preserving spread to U by setting $\rho = 5$, $\omega = 15$, which implies var$(U|I_{t-1}) = 8.333$.

Table 2

var$(U|I_{t-1})$ **and Endogenous Variables**

($M_I = 1$, $m_p = 1.1$, $\phi = 0.5$, $\psi = 0.1$, $h = 1$**)**

var$(U\|I_{t-1})$	q_l	$E(Y\|I_{t-1})$	var$(Y\|I_{t-1})$	$E(P\|I_{t-1})$	var$(P\|I_{t-1})$
5.3333	0.42	2.0412	0.0180	4.1928	0.4178
8.3333	0.83	2.0503	0.0699	4.1358	0.4083

As with quadratic hedging costs, the higher shock variance induces a higher fraction of learners, raising mean output and lowering the mean price. The higher variance of the shock lowers the variance of the price because learners' output responds more to demand shocks, dampening price fluctuations.

4 Learning and Hedging Under Monopolistic Competition

The New Keynesian model is the dominant framework for the analysis of monetary policy.[16] A key feature of this literature is the assumption that firms operate under monopolistic competition. By imposing some type of price rigidity, a New Keynesian Phillips Curve is generated which, along with an Euler Equation and interest rate rule, closes a simple version of the model. This section introduces the hedging and learning framework from the simple model of the previous section into a model of monopolistic competition. Unlike in much of the New Keynesian literature, we assume that prices are flexible and that firms choose their level of production. Most of the results from the cobweb model are preserved: hedging dampens the effect of demand shocks, a stable interior equilibrium typically exists, and the comparative statics behave in an intuitive manner. The introduction of an interest rate rule further reveals that our framework greatly complicates the conduct of monetary policy. Monetary policy now influences the steady state level of production by affecting the volatility of the economy and thus determining the fraction of firms that choose to learn. Optimal monetary policy requires that the monetary authority maximize steady state output, while placing less emphasis on stabilization.[17] Firms that learn produce more output than firms that hedge. Monetary policy that destabilizes prices now incentivizes more firms to learn and thus increases steady state output. This motivation implies

---

[16] For details on this model, see Woodford (2003) and Gali (2008).

[17] In standard versions of the New Keynesian Model, the monetary authority can influence the long run level of output. This ability, however, disappears as the discount factor approaches one and is therefore not typically a major emphasis of policy for realistic calibrations. See Clarida, Gali, and Gertler (1999) for further discussion.

that optimal policy ensures sufficiently volatile prices. Whereas monetary policy that stabilizes prices is optimal in most of the New Keynesian literature, such a policy performs poorly in our model.

Households

Our modeling of households is standard. A continuum over the unit interval of identical households maximize utility through consumption of a composite good (Y_t) and the supply of labor (N_t). They take the nominal wage (W_t) and price index (P_t) as given. We assume the following functional form of the representative agent's utility function.

$$Max_{Y_t,N_t} \;\; E_t[\sum_{i=0}^{\infty}\beta^i(\frac{ln(Y_{t+i})}{U_{t+i}} - \gamma N_{t+i})] \tag{25}$$

s.t.

$$b_t + P_tY_t = R_{t-1}b_{t-1} + N_tW_t + \Pi_t \tag{26}$$

U_t represents a white noise preference shock. Π_t represents aggregate firm profits which are returned to the representative household, which takes them as given. Households may purchase one-period bonds (b_t) from each other at the risk free interest rate ($R_t - 1$), but in equilibrium, $b_t = 0$. We assume that the monetary authority is able to set R_t at the level of its choosing.[18]

Optimization yields a standard Euler Equation and a first-order condition for labor supply.

$$\frac{1}{Y_tU_t} = E_t[\frac{\beta R_t}{\pi_{t+1}Y_{t+1}}] \tag{27}$$

$$\frac{W_t}{U_tY_tP_t} = \gamma \tag{28}$$

Firms

A continuum over the unit interval of ex-ante, identical firms exists. Each firm produces a differentiated good. Firms first choose whether to pay a learning cost that is proportional to the aggregate price level (ψP_t). By paying this cost, learning firms are able to perfectly observe U_t and therefore contemporaneously calculate the values of the model's other variables.[19] If a firm does not pay the

---

[18] This ability results from adding real money balances (M_t/P_t) to the utility function and generating a first-order condition for money demand. Because this equation is unimportant for the remainder of this section, we suppress it here.

[19] There are two equivalent interpretations of this type of learning. First, learning may equal the economic cost associated with measuring current market conditions. Second, U_t may be fully observable, but firms may have to choose their production with a one-period lag unless they pay the learning cost.

learning cost, and instead chooses to hedge, then it cannot observe U_t or correctly calculate the contemporaneous values of the model's other variables. Hedgers thus make their choices in period t based on an information set that only includes data through period $t-1$. Hedgers choose both their output and how much of their production to hedge (Z_t) at a cost equal to $hE_{t-1}[P_t]Z_t^2$.[20] Hedging firms are assured of receiving a strike price, $P_{s,t}$ for the hedged portion of their production.

Firms produce differentiated goods which are then costlessly transformed into a final composite good using the standard indices:

$$Y_t = (q_l Y_{l,t}^{\frac{\epsilon-1}{\epsilon}} + (1-q_l) Y_{h,t}^{\frac{\epsilon-1}{\epsilon}})^{\frac{\epsilon}{\epsilon-1}} \tag{29}$$

$$P_t = (q_l P_{l,t}^{1-\epsilon} + (1-q_l) P_{h,t}^{1-\epsilon})^{\frac{1}{1-\epsilon}} \tag{30}$$

where each type of firms' demand is:

$$P_{i,t} = P_t (\frac{Y_t}{Y_{i,t}})^{1/\epsilon} \tag{31}$$

We assume that both types of firms have mean-variance preferences. We further assume that for hedging firms, the strike price equals the firms' expected price conditional on not observing U_t, $E_{t-1}[P_{i,t}]$, and that U_t is uniformly distributed between $1-\frac{\eta}{2}$ and $1+\frac{\eta}{2}$ where $\eta < 2$. The representative hedging firm thus solves the following problem:

$$Max_{Y_{i,t},Z_t} \;\; E_{t-1}[\Pi_t^h - \varphi \frac{Var(\Pi_t^h)}{P_t}] \tag{32}$$

$$\Pi_t^h = Y_{i,t}^{\frac{\epsilon-1}{\epsilon}} P_t Y_t^{1/\epsilon} - cY_{i,t}W_t - hE_{t-1}[P_t]Z_t^2 \tag{33}$$

where $Y_{i,t} = \frac{N_{i,t}}{c}$ and φ represents the level of risk aversion. We assume that preferences depend on the expected "real" variance, $\frac{Var(\Pi_t^h)}{P_t}$. By relying on the real variance, we ensure that the model has the desirable property that no real variables depend on the steady state price level, which is typically indeterminate in the model.

The representative hedger's revenue may be broken down into the certain, hedged component $Z_t E_{t-1}[P_{i,t}]$, and the uncertain, non-hedged component, $(Y_{i,t} - Z_t)P_{i,t}$. The firm's wage bill, $cY_{i,t}W_t$ is uncertain, while the firm's hedging costs are known. Inserting Equations (28) and (31) into Equation (33) demonstrates that the variable component of the hedger's profits equals $Y_{i,t}^{-1/\epsilon}(Y_{i,t} - Z_t)P_t Y_t^{1/\epsilon} -$

---

[20] An alternate approach is to assume that the hedging costs depend on expected firm prices instead of the expected aggregate price level. Under this assumption, the model may exhibit multiple, stable steady states as in Section 2.

$c\gamma Y_{i,t}U_tP_tY_t$. Decomposing the variance allows us to rewrite the hedger's profit maximization problem.

$$Max_{Y_{i,t},Z_t} \;\; E_{t-1}[\frac{\Pi_t^h}{P_t} - \varphi Y_{i,t}^{-2/\epsilon}(Y_{i,t}-Z_t)^2\Omega_1 - \varphi c^2\gamma^2 Y_{i,t}^2\Omega_2 - 2\varphi c\gamma Y_{i,t}^{\frac{\epsilon-1}{\epsilon}}(Y_{i,t}-Z_t)\Omega_3] \tag{34}$$

where $\Omega_1 = E_{t-1}[\frac{Var(P_tY_t^{1/\epsilon})}{P_t^2}]$, $\Omega_2 = E_{t-1}[\frac{Var(P_tY_tU_t)}{P_t^2}]$, and $\Omega_3 = E_{t-1}[\frac{Covar(P_tY_t^{1/\epsilon},P_tY_tU_t)}{P_t^2}]$.

Differentiating with respect to $Y_{i,t}$ and Z_t, and using Equation (27) to eliminate Y_t yields two first-order conditions:

$$\tfrac{\epsilon-1}{\epsilon}Y_{h,t}^{\frac{-1}{\epsilon}}E_{t-1}[Y_t^{\frac{1}{\epsilon}}]^{1/\epsilon} - c\gamma E_{t-1}[Y_tU_t]...$$

$$-\varphi(\frac{2(\epsilon-1)}{\epsilon}Y_{h,t}^{\frac{\epsilon-2}{\epsilon}} - \frac{2(\epsilon-2)}{\epsilon}Y_{h,t}^{\frac{-2}{\epsilon}}Z_t - \frac{2}{\epsilon}Y_{h,t}^{\frac{-2-\epsilon}{\epsilon}}Z_t^2)\Omega_1 - 2\varphi c^2\gamma^2Y_{h,t}\Omega_2 + 2\varphi c\gamma(\frac{2\epsilon-1}{\epsilon}Y_{h,t}^{\frac{\epsilon-1}{\epsilon}} - \frac{\epsilon-1}{\epsilon}Y_{h,t}^{1/\epsilon}Z_t)\Omega_3 = 0 \tag{35}$$

$$-hZ_t + \varphi(Y_{h,t}^{\frac{\epsilon-2}{\epsilon}} + Z_tY_{h,t}^{\frac{-2}{\epsilon}})\Omega_1 + \varphi c\gamma Y_{h,t}^{\frac{\epsilon-1}{\epsilon}}\Omega_3 = 0 \tag{36}$$

The representative learning firm's problem is simpler. Learning eliminates all uncertainty and the variance of profits may thus be ignored in the profit maximization problem. We further assume that the representative learning firm may not hedge any of its output.[21]

$$Max_{Y_{i,t}} \;\; Y_{i,t}^{\frac{\epsilon-1}{\epsilon}}P_tY_t^{1/\epsilon} - \psi P_t - cY_{i,t}W_t \tag{37}$$

Which yields the following first-order condition:

$$Y_{l,t} = (\frac{\epsilon-1}{\epsilon c\gamma})^{\epsilon}Y_t^{1-\epsilon}U_t^{-\epsilon} \tag{38}$$

Finally, the fraction of firms that learn, q_l, is determined by setting equal the expected mean-variance profits of learning and hedging. Any interior value of q_l satisfies:

$$E_{t-1}[\frac{\Pi_t^h}{P_t} - \varphi Y_{h,t}^{-2/\epsilon}(Y_{h,t}-Z_t)^2\Omega_1 - c^2\gamma^2Y_{h,t}^2\Omega_2 - 2c\gamma Y_{h,t}^{\frac{\epsilon-1}{\epsilon}}(Y_{h,t}-Z_t)\Omega_3)] = E_{t-1}[\Pi_t^l] \tag{39}$$

---

[21] Allowing learning firms to hedge, prior to observing U_t, would not affect our results. Because learning eliminates all uncertainty, learning firms would choose zero hedging for any $h > 0$.

Hedging and Learning Suppliers

We now consider the details of how firms learn and hedge. For now, we take the simplest approach.[22] We assume that a learning supply sector costlessly acquires information on behalf of learning firms and thus realizes profits equal to $q_l \varphi P_t$, which are then returned to the representative household. Likewise a hedging supply sector purchases hedgers' output at the strike price and costlessly re-sells this output to households at the realized price. The resulting profits, $(1-q_l)Z_t(E_{t-1}[P_{h,t}] - P_{h,t} + hZ_tE_{t-1}[P_t])$, are then returned to the representative household.

Equilibrium

We define an equilibrium as any sequence of $q_l, Y_{l,t}, Y_{h,t}, P_{l,t}, P_{h,t}, Z_t, Y_t$ and P_t such that Equations (27), (29)-(31), (35)-(36), and (38)-(39) are satisfied. We calculate equilibrium by relying on a partial first-order log-linearization around a zero inflation steady state where $\pi_t = \frac{P_t}{P_{t-1}} = 1$.

Log-linearizing the Euler Equation, (27):

$$\tilde{Y}_t = E_t[\tilde{Y}_{t+1} + \tilde{\pi}_{t+1}] - \tilde{R}_t - \tilde{u}_t \tag{40}$$

Log-linearizing learning firm's first order condition, (38):

$$\bar{Y}_l \tilde{Y}_{l,t} = (1-\epsilon)(\frac{\epsilon - 1}{c\epsilon\gamma})^{\epsilon} \bar{Y}_t^{1-\epsilon} \tilde{Y}_t - \epsilon(\frac{\epsilon - 1}{c\epsilon\gamma})^{\epsilon} \bar{Y}_t^{1-\epsilon} \tilde{u}_t \tag{41}$$

Log-linearizing the output index, (29):

$$\bar{Y}_t^{\frac{\epsilon-1}{\epsilon}} \tilde{Y}_t = q_l \bar{Y}_{l,t}^{\frac{\epsilon-1}{\epsilon}} \tilde{Y}_{l,t} + (1-q_l) \bar{Y}_{h,t}^{\frac{\epsilon-1}{\epsilon}} \tilde{Y}_{h,t} \tag{42}$$

where $\bar{X}$ is the steady state value of X_t and $\tilde{X}_t = \frac{X_t - \bar{X}}{\bar{X}}$.

The lack of serial correlation in our model greatly simplifies the analysis by allowing us to treat $Y_{h,t}$, Z_t, Ω_1, Ω_2, Ω_3, and q_l as constants. We therefore do not need to log-linearize the first order conditions for hedging firms.

A Simple Case, Constant Nominal Interest Rate

To close the model, we specify a rule by which the monetary authority sets R_t. In most of the New Keynesian literature, the monetary authority's problem is choosing the rule that best stabilizes the economy around an exogenous steady state. Our assumption that firms are risk averse, however, greatly complicates monetary policy. By affecting the model's stability, monetary policy influences

---

[22] Section 5 considers the implications of endogenizing both learning and hedging costs.

the incentives to learn and hedge and thus steady state output. We begin, however, by analyzing a simple case where the nominal interest rate is constant so that $R_t = \beta^{-1}$.

Because U_t is white noise, the model exhibits no serial correlation, $Y_{h,t} = \bar{Y}_h$, and $Z_t = \bar{Z}\ \forall t$. Additionally, q_l is a constant. Hedging firms are thus completely unresponsive to preference shocks. Combining Equations (41) and (42) and using the method of undetermined coefficients yields a log-linearized expression for output.[23]

$$\tilde{Y}_t = \frac{q_l(1-\epsilon)(\frac{\epsilon-1}{\epsilon c\gamma})^{\epsilon}\bar{Y}_l^{-1/\epsilon}\bar{Y}^{1-\epsilon}}{\bar{Y}^{\frac{\epsilon-1}{\epsilon}} - q_l(1-\epsilon)(\frac{\epsilon-1}{\epsilon c\gamma})^{\epsilon}\bar{Y}_l^{-1/\epsilon}\bar{Y}^{1-\epsilon}}\tilde{u}_t = \tau\tilde{u}_t \tag{43}$$

Restating Equation (40) and taking the informed expectation:

$$\tilde{Y}_t = -\tilde{P}_t - \tilde{u}_t \tag{44}$$

$$\tilde{P}_t = (-\tau - 1)\tilde{u}_t \tag{45}$$

We now use Equations (43) and (45) to calculate closed form solutions that approximate Ω_1, Ω_2, and Ω_3 for any value of q_l, and thus allow us to solve for the model's steady state, the values of $\alpha_1 - \alpha_6$, and hence τ. Using Equations (44) and (45) to compute the variance of $P_t Y_t U_t$:

$$\Omega_2 = \bar{Y}^2 \frac{\tau^2(\tau+1)^2(\sigma_u^2)^3 + (\tau^2+\tau+1)(\sigma_u^2)^2}{(\tau+1)^2\sigma_u^2 + 1} \tag{46}$$

Tedious, but straightforward integration by parts solves for Ω_1 and Ω_3 as more complicated functions of ϵ, η, τ, and $\bar{Y}$.

Equation (43) also allows us to compute the constant expectations that appear in the representative hedging firm's first order condition with respect to $Y_{h,t}$, (35).

$$E_{t-1}[Y_t U_t] = \bar{Y}(1 + \frac{\tau\eta^2}{4}) \tag{47}$$

$$E_{t-1}[Y_t^{\frac{1}{\epsilon}}] = \frac{\epsilon\bar{Y}^{\frac{1}{\epsilon}}}{(\epsilon+1)\eta\tau}[(1+\frac{\eta\tau}{2})^{\frac{\epsilon+1}{\epsilon}} - (1-\frac{\eta\tau}{2})^{\frac{\epsilon+1}{\epsilon}}] \tag{48}$$

---

[23] The solution represented by Equation (43) is not necessarily unique, but instead represents the model's minimum state variable solution. A major objective of monetary policy is to ensure a unique equilibrium that does not depend on extraneous factors (sunspots). We rule sunspot equilibria out *a priori*. Ordinarily, sunspot equilibria are undesirable because they introduce additional volatility into the system. In our model, however, sunspot equilibria may be desirable because the additional volatility may induce more learning and hence increase steady state output. We leave further examination of sunspot equilibria for future research.

Figure 3: Expected Learning Less Hedging Profits

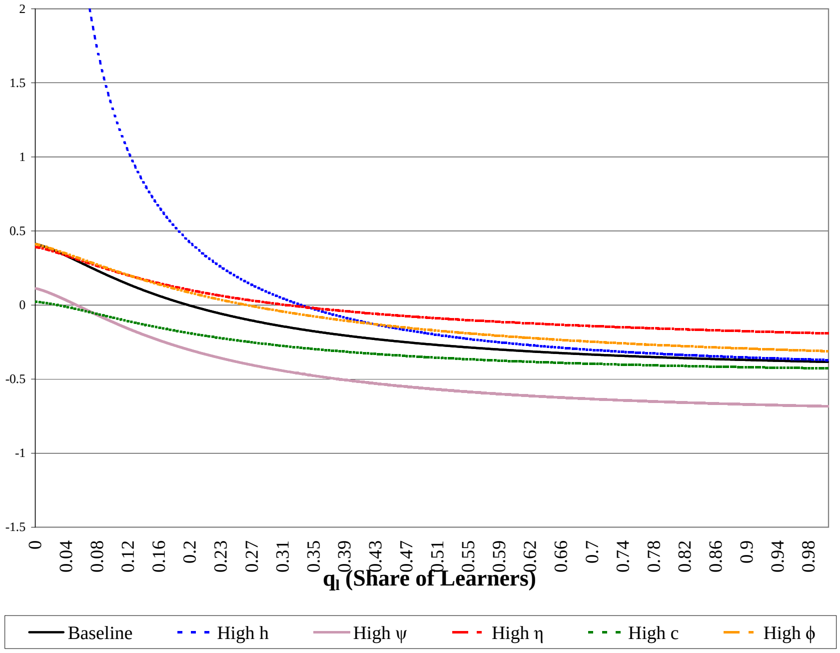

Calibration and Simulation

Our baseline calibration sets $\beta = 0.99$, $\epsilon = 20$, $\varphi = 0.5$, $c = 0.7$, $\gamma = 1$, $\psi = 0.5$, $\eta = 0.1$, and $h = 0.01$. We also consider calibrations with increased risk aversion ($\varphi = 1.0$), decreased productivity ($c = 0.9$), increased learning costs ($\psi = 0.8$), more volatile preference shocks ($\eta = 0.3$), and increased hedging costs ($h = 0.1$).

Figure 3 displays the expected profits under learning less those under hedging for each q_l.

A common pattern exists across all calibrations. The volatility of profits closely tracks that of the aggregate price level. If q_l is low, then preference shocks are passed on to the price level instead of output, and firm level prices are highly volatile. Learning is thus preferable to hedging. As q_l increases, preference shocks are increasingly passed on to output instead of prices, and learning becomes less appealing. Under the baseline calibration, a unique and stable steady state exists when $q_l = 19.25\%$.

Increased risk aversion ($\varphi = 1.0$) increases the incentives to both learn and hedge and the steady state now occurs where $q_l = 26.65\%$.[24] Lower productivity ($c = 0.9$) increases the learning cost as a share of profits and reduces the steady state share of learning firms to $q_l = 2.85\%$. Higher learning costs ($\psi = 0.8$) displace the baseline calibration downwards. The steady state now occurs where

---

[24] This result contrasts with that from the cobweb model where heightened risk aversion decreased learning.

Figure 4: Fraction of Output that Hedging Firms Hedge

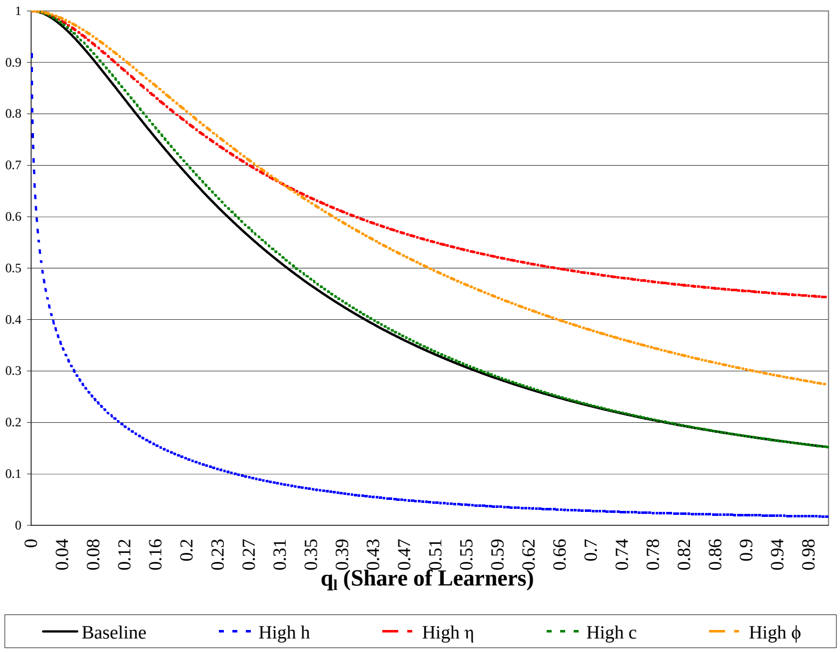

$q_l = 5.15\%$. Increasing the support of preference shocks ($\eta = 0.3$) has similar effects as increasing risk aversion. The steady state now occurs where $q_l = 31.65\%$. Finally, increased hedging costs ($h = 0.1$) make learning more attractive and the unique steady state now occurs where $q_l = 33.25\%$.[25]

Because hedging firms face uncertainty, they produce less output than learning firms. Aggregate steady state output therefore always increases along with q_l. As q_l increases, both types of steady state profits decline, although greater price stabilization benefits hedging firms and causes their (mean-variance) profits to fall at a slower rate.

Figure 4 reports the fraction of output that hedging firms choose to hedge.

The results are similar for each calibration.[26] As q_l increases, more firms observe preference shocks and the response of aggregate output is thus larger. The effect on the price level is inversely related to τ, thus more learners leads to more price stability. Hedging firms thus choose to hedge a smaller share of their production. Unsurprisingly, a higher hedging cost reduces the share of production that is hedged while higher risk aversion and larger preference shocks have the opposite effect.

Figure 5 reports τ, the responsiveness of output to preference shocks.

---

[25] As in the simpler model of the previous sections, however, very large hedging costs typically do not induce all firms to learn. Instead, some firms will choose to not learn and hedge only a small fraction of their output.

[26] Figures 4-4 do not display the results for a heightened learning cost ($\psi = 0.8$) because they are identical to the baseline case.

Figure 5: τ (Response of Output to Preference Shocks)

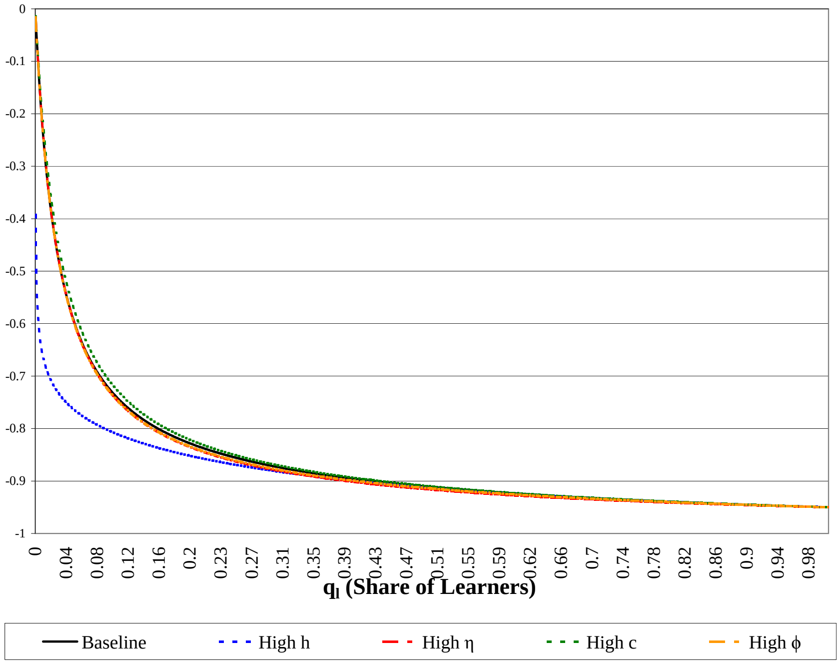

All calibrations behave in a similar manner. As q_l increases, more firms respond to preference shocks and the magnitude of τ increases. The responsiveness of the price level is inversely related to τ and thus decreases with q_l. Table 3 reports the steady state properties for each calibration.[27]

Table 3

Steady State Properties for Alternate Calibrations

Calibration	q_l	$\bar{Y}$	$\bar{Y}_h$	$\bar{Y}_l$	$\frac{\bar{Z}_t}{\bar{Y}_h}$	τ	Ω_1
high c	0.0285	1.0340	1.0188	1.5609	0.9865	-0.4447	1.4188
high ψ	0.0515	1.3247	1.2805	2.1492	0.9521	-0.6078	0.3950
baseline	0.1925	1.3383	1.2363	1.7700	0.6873	-0.8267	0.0427
high ϕ	0.2665	1.3398	1.1988	1.7317	0.7179	-0.8660	0.0234
high η	0.3175	1.3410	1.1741	1.7032	0.6619	-0.8833	0.0204
high h	0.3355	1.3417	1.1690	1.6870	0.0746	-0.8879	0.0156

These results allow us to examine the desirability of each steady state. As q_l increases, aggregate output typically increases and is always maximized when $q_l = 99.95\%$.[28] If welfare is based on

---

[27] Revenue uncertainty is the major source of hedgers' risk. Table 3 thus reports Ω_1 while omitting Ω_2 and Ω_3.

[28] We assume that at least 0.05% of firms always choose to learn and that an equal fraction always choose to hedge.

the utility of the representative household, then higher levels of q_l increase steady state welfare.[29] Although the monetary authority must also consider the effects of its policy on volatility, the next section demonstrates that the impact on steady state output dominates this latter concern.

Lower productivity (high c) alone affects the production function and this calibration unsurprisingly results in the lowest level of steady state output. Additionally, the relative size of the learning cost is especially high and as a result, q_l is low. For the remaining calibrations, output (and expected utility) are increasing in q_l. Higher risk aversion (high φ) and larger support for the stochastic shocks provide incentives for both more hedging and more learning. In both cases, the resulting increase in learning is sufficient to increase steady state output, even though firms that hedge produce less. Higher learning costs (high ψ) clearly deters learning and results in reduced steady state output relative to the baseline. Finally, higher hedging costs (high h) induce firms to substitute away from hedging to learning and results in the highest level of steady state output. This result suggests that taxes or other policies that discourage hedging may be welfare improving.

Monetary Policy

The results for $R_t = \beta^{-1}$ demonstrate that the monetary authority faces an unusually complicated problem. Rather than simply attempting to affect the volatility of output, by affecting volatility, monetary policy influences the steady state level of output through supply side effects. We now examine monetary policy under commitment by assuming the following interest rate rule.

$$\tilde{R}_t = \lambda_u \tilde{u}_t + \tilde{e}_t \tag{49}$$

where e_t is a white-noise control error which, for tractability, we assume is uniformly distributed between $-\frac{\delta}{2}$ and $\frac{\delta}{2}$. We continue to assume that hedging firms must choose $Y_{h,t}$ based on $t-1$ information and thus cannot observe R_t.

The only change to the model is in the log-linearized Euler Equation, (40).

$$\tilde{Y}_t = E_t[\tilde{Y}_{t+1} + \tilde{\pi}_{t+1}] - \lambda_u \tilde{u}_t - \tilde{e}_t - \tilde{u}_t \tag{50}$$

We obtain the solution for the price level by again taking the informed expectation and exploiting the lack of serial correlation in the model:

$$\tilde{P}_t = -(\tau + \lambda_u + 1)\tilde{u}_t - \tilde{e}_t \tag{51}$$

---

[29] Increasing $\bar{Y}$ increases both the utility from consumption and the disutility of labor, and could thus potentially reduce steady state welfare. The assumption of monopolistic competition, however, introduces a distortion into the model that causes firms to inefficiently underproduce. Steady state utility is thus increasing in steady state output in the model.

Ω_1, Ω_2, and Ω_3 may again be approximated as complex functions of the model's structural parameters. Because monetary policy affects both the first and second moments of household utility, it is not appropriate to evaluate monetary policy rules by comparing the weighted variances of inflation and output as in much of the New Keynesian literature. Instead, we integrate the representative household's utility function, (25), over the distribution of $\tilde{u}_t$.

$$\int_{-\eta/2}^{\eta/2}[\tfrac{ln(Y_t)}{U_t} - c\gamma Y_t]\partial\tilde{u}_t =$$

$$[ln(\tau\tilde{u}_t + 1)ln(\bar{Y}) + ln(\tau\tilde{u}_t + 1)ln(1 - \frac{\tau\tilde{u}_t + 1}{1-\tau}) + Li_2(\frac{\tau\tilde{u}_t + 1}{1-\tau}) - \frac{c\bar{Y}\gamma\tau\tilde{u}_t^2}{2}]|_{-\eta/2}^{\eta/2} - c\bar{Y}\gamma\tau \quad (52)$$

where Li_2 indicates a second-order polylogarithmic function.

We now simulate the high hedging cost calibration for λ_u ranging between -0.5 and 2. We set the monetary authority's control error equal to zero ($\delta = .00$).[30] In much of the New Keynesian model, optimal monetary policy eliminates the distortion resulting from sticky prices by achieving price stabilization. Our model has a similar stabilization motive. Monetary policy ideally allows output to efficiently respond to preference shocks rather than allowing the entire impact to be passed on to the price level. More importantly, however, excessive price stability reduces the incentive to learn and fewer learners result in less steady state output. Optimal monetary policy thus allows for an intermediate level of price volatility which induces additional learning. Figure 6 plots both the steady state q_l (left vertical axis) and a measure of aggregate price volatility, $(\tau + \lambda_u + 1)^2$ (right vertical axis).

As λ_u increases, so does the volatility of the price level. This increases the uncertainty of hedging firms and induces a higher amount of learning. If $\lambda_u \geq 0.12$, then all firms choose to learn and output is maximized. The importance of the first moment of utility ensures that it is optimal to pursue a policy that results in $q_l = 99.95\%$. As λ_u increases above 0.12, however, prices are further destabilized causing the small fraction of remaining hedgers (0.05%) to reduce their output.[31] Utility is thus maximized when $\lambda_u = 0.12$.

The study of rules-based monetary policy usually assumes that the interest rate rule targets measures of prices and output instead of fundamental shocks. The optimal interest rate rule from Equation (49), however, is isomorphic to a rule targeting only expected inflation.

$$\tilde{R}_t = \lambda_\pi E_t[\tilde{\pi}_{t+1}] + \tilde{e}_t \quad (53)$$

---

[30] The results are qualitatively unchanged for positive values of δ.

[31] If all firms are allowed to learn, then any policy that achieves q_l=1 is optimal.

Figure 6: Steady State Under Different Policy Rules (δ=0)

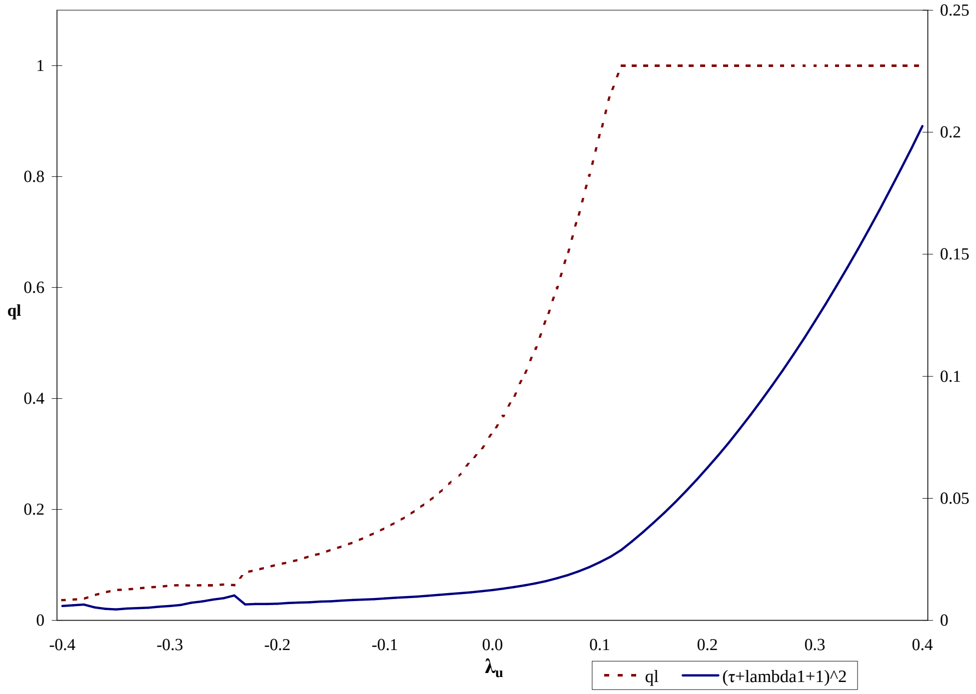

Substituting Equation (53) into the linearized Euler Equation, (50) results in an expression for the price level that depends on λ_π instead of λ_u. Setting this expression equal to Equation (51) converts the interest rate rule into an expected inflation target.

$$\lambda_\pi = -[\frac{\tau + 1}{\tau + 1 + \lambda_u} - 1] \tag{54}$$

For $\lambda_u = 0.12$, the equivalent expected inflation target sets $\lambda_\pi = 0.71$.[32]

For the previous results, the monetary authority's stabilization motive and its desire to maximize steady state output are compatible. By maximizing learning, optimal policy also maximizes firms' ability to effectively respond to preference shocks. To demonstrate the dominance of supply side effects, however, we now vary λ_1 between -2 and 2 for a constant fraction of learners ($q_l = 33.35\%$) using the high hedging cost calibration. Figure 7 reports the results.

At $\lambda_u = -0.13$, price volatility is minimized, and output and utility are both maximized. By stabilizing prices, monetary policy reduces the uncertainty that hedging firms face and thus increases both their output and aggregate output. If q_l is exogenous, then such a policy is now optimal. In this case, the monetary authority is not able to satisfy its stabilization motive. Highly stable prices increase

---

[32] The presence of only one shock implies that the monetary authority needs only one instrument to implement the optimal policy. A continuum of equivalent rules that target both output and inflation therefore also exists.

Figure 7: Effect of Policy on Welfare and Price Volatility, $q_l = 33.35\%$

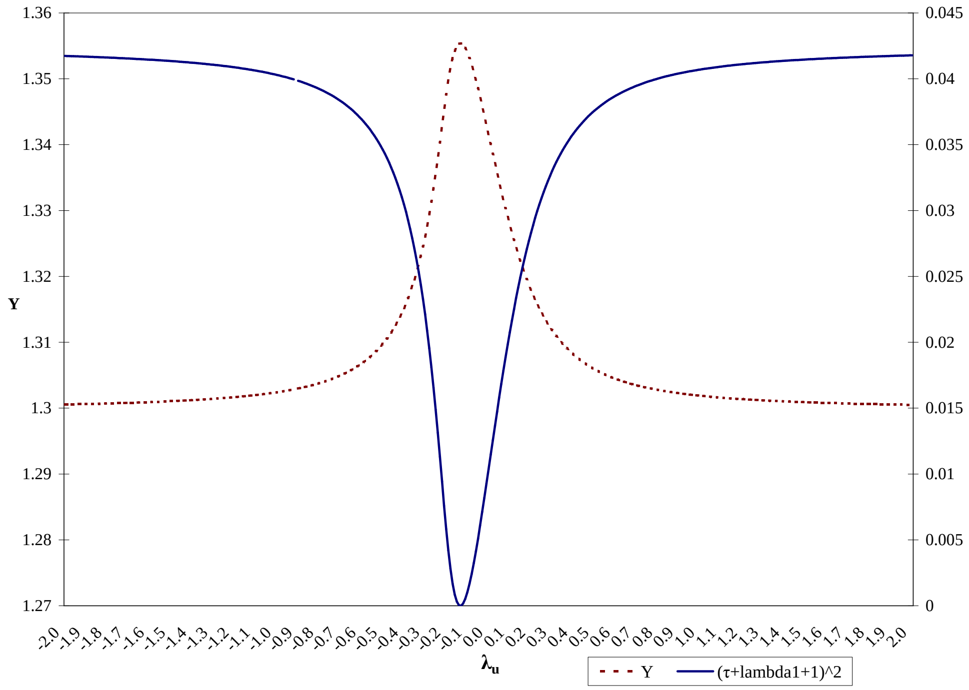

the share of output produced by hedgers, inefficiently dampening output's response to preference shocks.[33] The effect on output's first moment, however, dominates the effect on its second moment.

5 Endogenous Learning and Hedging

The previous section treats the costs of learning and hedging, ψ, and h, as exogenous parameters. Furthermore, the analysis assumes that there is no social cost to either hedging or learning. This section develops microfoundations for both types of risk management by assuming that the learning and hedging sectors must employ labor to perform these activities. Because the representative household obtains disutility from supplying this additional labor, this approach potentially affects our result that optimal monetary policy maximizes both steady state output and employment.

Learning

We define κ as the constant labor requirement for a learning firm to acquire information regarding its demand shock. The representative learning firm's profit maximization problem thus becomes:

---

[33] At $\lambda_u = -0.13$, τ is maximized at -0.86. Its most efficient value equals -0.95 when all firms learn.

$$Max_{Y_{i,t}} \quad Y_{i,t}^{\frac{\epsilon-1}{\epsilon}} P_t Y_t^{1/\epsilon} - W_t\kappa - cW_tY_{i,t} \tag{55}$$

The resulting first order condition is unaffected by this modification.

$$Y_{l,t} = (\frac{\epsilon-1}{c\epsilon\gamma})^{\epsilon} Y_t^{1-\epsilon} U_t^{-\epsilon} \tag{56}$$

Hedging Suppliers

We assume the existence of a competitive hedging sector with free entry and exit. To be consistent with a competitive hedging sector, we drop our assumption that a hedging firm pays hedging costs equal to $hZ_t^2E_{t-1}[P_t]$, and instead assume that hedging costs simply equal hZ_t. We assume that each hedging supplier incurs a fixed cost, ιW_t, and must pay variable labor costs that are proportional to the amount of production that is hedged, $\upsilon Z_{i,t}W_t$. We assume that hedging suppliers have the same mean-variance preferences as firms. Like producers that hedge, hedging suppliers face risk from both the wage rate and the equilibrium price of hedgers' output. The optimization problem for a hedging supplier thus becomes:

$$Max_{Z_{i,t}} \quad E_{t-1}[hZ_{i,t} - \iota W_t - \upsilon Z_{i,t}W_t ...$$

$$-\varphi P_t^{-1}(Z_{i,t}^2 Var(P_{h,t}) + \varphi(\upsilon^2 Z_{i,t}^2 + \iota^2)Var(W_t) + 2\varphi Z_{i,t}(\upsilon Z_{i,t} + \iota)Cov(P_{h,t}, W_t))] \tag{57}$$

Using Equations (28) and (31), the representative hedging supplier's problem may be re-written in terms of Ω_1, Ω_2, and Ω_3.[34]

$$Max_{Z_{i,t}} \quad E_{t-1}[\frac{hZ_{i,t}}{P_t} - \gamma(\iota+\upsilon Z_{i,t})Y_tU_t - \varphi Y_{h,t}^{\frac{-2}{\epsilon}}(Z_{i,t}^2\Omega_1 + \gamma^2(\iota+\upsilon Z_{i,t})^2\Omega_2 + \gamma(\iota+\upsilon Z_{i,t})Y_{h,t}^{\frac{-1}{\epsilon}}Z_{i,t}\Omega_3)]] \tag{58}$$

Differentiating with respect to $Z_{i,t}$ and eliminating W_t using (28) yields a first order condition for the representative hedging supplier:

$$Z_{i,t} = E_{t-1}[\frac{\frac{h}{P_t} - \gamma\upsilon Y_tU_t - \varphi\gamma\iota(2\gamma\upsilon\Omega_2 + Y_{h,t}^{\frac{-1}{\epsilon}})}{2\varphi(Y_{h,t}^{\frac{-2}{\epsilon}}\Omega_1 + \gamma^2\upsilon^2\Omega_2 + \gamma\upsilon Y_{h,t}^{\frac{-1}{\epsilon}}\Omega_3)}] \tag{59}$$

We now impose a zero-profits condition to determine the equilibrium value of $Z_{i,t}$ h, and corresponding number of hedging suppliers.

---

[34] We treat the output of hedging firms, $Y_{h,t}$ as observable to hedging suppliers.

$$E_{t-1}[\tfrac{h}{P_t}] = E_{t-1}[\gamma\upsilon Y_t U_t + 2\varphi\gamma^2\iota\upsilon\Omega_2 + \varphi\gamma\iota Y_{h,t}^{\frac{-1}{\epsilon}}\Omega_3 ...$$

$$+2\sqrt{\varphi(\iota\upsilon Y_t U_t + \varphi\gamma^2\iota^2\Omega_2)(Y_{h,t}^{\frac{-2}{\epsilon}}\Omega_1 + \gamma^2\upsilon^2\Omega_2 + \gamma\upsilon Y_{h,t}^{\frac{-1}{\epsilon}}\Omega_3)}] \tag{60}$$

Hedging Firms

The hedging firm's first order condition with respect to $Y_{h,t}$, (35), is unaffected. The representative hedging firm takes h as given, and the first order condition with respect to Z_t now becomes:

$$\frac{h}{2} - \varphi\Omega_1(-Y_{h,t}^{\frac{\epsilon-2}{\epsilon}} + Y_{h,t}^{\frac{-2}{\epsilon}} Z_t) - \varphi c\gamma Y_{h,t}^{\frac{\epsilon-1}{\epsilon}}\Omega_3 = 0 \tag{61}$$

Households

The representative household's instantaneous utility now equals $\frac{ln(Y_t)}{U_t} - \gamma(N_t + (1-q_l)(\iota\frac{Z_t}{Z_{i,t}+\upsilon Z_t} + q_l\kappa)$. The households first order conditions, (27) and (28) are unaffected. The extra labor supply for hedging and learning does, however, affect the optimality of monetary policy.

The model is then solved in the same manner as in the previous section. Its behavior is also similar. As more firms learn, uncertainty is reduced and steady state output increases. Monetary policy, however, now faces an additional tradeoff. Monetary policy that allows for sufficient volatility continues to maximize learning and steady state output, but this also has effects on the representative household's disutility of labor; more learning increases labor supply while to corresponding decline in hedging decreases labor supply. If learning costs are large, relative to hedging costs, then optimal monetary policy may entail stabilizing prices, minimizing learning, and reducing steady state output. For smaller learning costs, optimal monetary policy continues to induce maximum learning which maximizes steady state output.

6 Conclusion

The results from both models demonstrate that allowing risk-averse firms to actively manage their risk has serious and surprising implications. In both models, the steady state level of output (and expected utility in the latter model) closely tracks the fraction of firms that choose to learn. Parameter changes that induce more learning, such as higher hedging costs, lower learning costs, and larger stochastic shocks therefore increase output's natural rate.[35] The implications for monetary policy are striking and surprising. Policy makers should focus on output's first moment, not its second, and should allow sufficient price volatility to induce the maximum level of learning.

This paper analyzes the effects of allowing risk-averse firms to manage risk in a relatively simple monetary model. It would be of interest to examine how the inclusion of additional features, common to monetary models, would affect our results. Most New Keynesian models, for example, include serially correlated productivity shocks, as well as nominal rigidities. Likewise, for simplicity, we assume that learning eliminates all uncertainty. It would also be of interest to instead assume that learning firms receive only a noisy signal related to the model's stochastic shocks. Extending our framework to allow for these features is left for future research.

The related literature includes an enormous number of extensions of the New Keynesian framework. Given the significance of our results for Central Bankers, it is worthwhile to extend them to many of these. We list a few of the most promising avenues: First, an extensive literature examines how monetary policy may ensure the existence of a unique equilibrium. In this paper, we suppress this issue by ruling out sunspot equilibria. Whereas sunspot equilibria are typically welfare reducing, in our model they may be desirable as a way to increase uncertainty and thus increase steady state output. Second, we assume rational expectations where agent's know the model's reduced form solution. The adaptive learning literature instead assumes that agents form expectation using standard econometric techniques.[36] Adaptive learning introduces additional volatility into a model and it is of interest to examine how this extra uncertainty affects risk averse agents.[37] Finally, our model offers several novel opportunities for taxation to affect welfare.[38] Our results suggest that subsidizing firms' learning costs or taxing hedging may increase steady state output. These issues are left for future research.

---

[35] Increased risk aversion increases hedging and reduces learning in the cobweb model but has the opposite effect in the model of monopolistic competition.

[36] Adaptive learning is distinct from the learning discussed in this paper, which refers to the acquisition of knowledge about the model's stochastic shocks.

[37] For a general discussion of adaptive learning, see Evans and Honkapohja (2001). For an application of adaptive learning to monetary policy, see Orphanides and Williams (2007).

[38] The New Keynesian literature often advocates an optimal, if unrealistic, subsidy of employment to correct the distortion introduced by monopolistic competition. See Woodford (2003).

References

Amihud, Y. and B. Lev. 1981. "Risk Reduction as Managerial Motive for Conglomerate Mergers." *Bell Journal of Economics*, 12: 605-617.

Christiano, L., Eichenbaum, M., and C. Evans. 1997. "Sticky Price and Limited Participation Models of Money: a Comparison." *European Economic Review*, 41: 1201-1249.

Clarida, R., Gali, J., and M. Gertler. 1999. "The Science of Monetary Policy: A New Keynesian Perspective." *Journal of Economic Literature*, 37: 1661-1707.

Cooley, T., and K. Nam. 1998. "Asymmetric Information, Financial Intermediation and Business Cycles." *Economic Theory*, 12: 599-620.

Evans, G, and S. Honkapohja. 2001. *Learning and Expectations in Macroeconomics*. Princeton, NJ: Princeton University Press.

Gali, J. 2008. *Monetary Policy, Inflation, and the Business Cycle*. Princeton. Princeton University Press.

Geczy, C., Minton, B., and C. Schrand. 1997. "Why Firms Use Currency Derivatives." *Journal of Finance*, 52: 1323-1254.

Lucas, R. 1987. *Models of Business Cycles*. New York: Basil Blackwell.

Mankiw, G. and R. Reis. 2004. "Sticky Information versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve." *Quarterly Journal of Economics*, 117(4): 1295-1328.

May, D. 1995. "Do Managerial Motives Influence Firm Risk Reduction Strategies." *Journal of Finance*, 50(4): 1291-1308.

Orphanides, A. and J. Williams. 2007. "Robust Monetary Policy with Imperfect Knowledge." *Journal of Monetary Economics*, 54: 1407-1435.

Tufano, P., 1996. "Who manages risk? An empirical examination of risk management practices in the gold mining industry." *Journal of Finance*, 51: 1097-1137.

Taylor, J. 1980. "Aggregate Dynamics and Staggered Contracts." *Journal of Political Economy*, 88(1): 1-23.

The Wharton School and the Chase Manhattan Bank, N. A. 1995. "Survey of Derivative Usage Among U.S. Non-Financial Firms." Executive Summary.

Woodford, M. 2003. *Interest and Prices*. Princeton, NJ and Oxford: Princeton University Press.